Chapter 3

Establishing our Framework

Our interest in exploring presentation space arose from the intersection of two problems: one, a computer's display space is extremely limited in comparison to the amount of information it can store, and two, as the computer is an interactive medium, it is often not enough to create a static display when visualizing information.

The screen real estate problem is significant and likely to remain so for a while. The necessity for effective solutions to this problem is intensifying as the ability to produce visual data in greater volumes continues to outstrip the rate at which display technology is developing. We address the development of a better understanding of how utilizing presentation space can aid the screen real estate problem.

Currently, the space of presentation possibilities is relatively unexplored (see Chapter 1 and Chapter 2, Section 2.5). This is largely because the existence of an interactive presentation space is unique to the computer, and the computer has only recently been considered as an information medium. To visualize information it is no longer enough to create a static display of the information. While use of presentation in other mediums has focused on such things as gaining attention (in advertising), providing a perspective (in documentaries) and setting the mood (in theatre), most research into computational presentation space has concentrated on supporting exploration of unknown or partially known aspects of an information representation and more effective use of display space.

As described in Chapter 2, previous research into effective screen space usage has resulted in the development of various presentation methods. However, other than windows these methods have not been widely accepted. This is in spite of the fact that several readability issues with respect to windows have been identified, and that the new methods have been created to address them. Each new method appears distinct in functionality, appearance and the readability advantages offered. Also, they have often been tailored to particular information characteristics and/or tasks. As a result, at this point when developing a new interactive visualization, choices have to be made as to which readability issues are most important. These choices are sometimes limited by the information representation's characteristics.

In Chapter 1, we noted that research in information visualization had progressed to the point where interest in developing a general understanding of both the process of information visualization and the space of possible information visualizations was increasing [19, 30, 135, 165]. A framework that provides sufficient structure to enable understanding of the relationships between the individual known presentation methods and promotes insight into computational presentation space in general would be useful. There are two research contributions that provide frameworks for particular aspects of presentation space: Leung and Apperley's mathematical framework [94] and Furnas and Bederson's Space Scale Diagrams [53].

Leung and Apperley's [94] mathematical framework describes one way of relating distortion methods (Chapter 2, Sections 2.2.1 and 2.2.4). However for some methods, notably Perspective Wall [99], the resulting functions are not simple or intuitive. Furthermore distortion approaches are just a subset of presentation approaches. This subset is limiting in that all distortion approaches appear to eliminate some advantages of windowing, such as freedom of re-positioning, and those of ZUIs, such as the freedoms that come with extreme magnification.

Space Scale Diagrams [53] as described by Furnas and Bederson are a visual representation of the zooming relationships in ZUIs. Space Scale Diagrams have supported many insights including the fact that in a zooming environment the shortest path involves zooming out almost to the point where both ends of the trip are visible and then zooming back in. This path is shortest in terms of computational effort (the amount of image that has to be redrawn) and human cognitive effort. Users found it easiest to zoom out, change focus and zoom in, in that they were much more certain of the location in computational space during the process. They also noted that Space Scale Diagrams can describe fisheye, or distortion presentations. However, it is not clear how Space Scale Diagrams could be extended to describe multi-scale distortions with multiple foci.

This chapter proceeds as follows. Section 3.1 explains the decisions taken in establishing *Elastic Presentation Space* (EPS). Section 3.2 lists the desirable detail-in-context features. The next sections use EPS to develop a detail-in-context method that includes these features. Section 3.3 describes the creation of a single focus in the centre of the field of view. Section 3.4 describes the provision of choice of arbitrary focal size and shape. Section 3.5 discusses the control of the magnification of the focus. Section 3.6 describes freedom of focal location and how uniform magnification response for all focal positions is ensured. Section 3.7 describes the inclusion of multiple foci. Section 3.8 provides a brief discussion and Section 3.9 outlines how the usefulness of EPS will be verified.

As this framework is dependent on perspective geometry, Appendix B provides basic terminology and a brief review of perspective geometry.

3.1 EPS: The Basic Concepts

We present a framework, *Elastic Presentation Space* (EPS), that encompasses the presentation distortion dimension as defined in Chapter 2, Section 2.5. It incorporates insets, full zooming environments as described in Space Scale Diagrams [53], Leung and Apperley's distortion classifications [94] and more complex distortions that allow focal repositioning [25]. *Elastic* is a positive word that implies adjusting shape in resilient manner. That is, elastic materials can always revert to their original shape with ease. This pertains to the Piagetian notions of reversibility and revertability [29, 121] and how necessary being able to return to previous configurations is to create a environment that has sufficient security for intellectual exploration.

The basic concept of EPS is:

- a two-dimensional visual representation is placed onto a surface;
- this surface is placed in three-dimensional space;
- the surface, containing the representation, is viewed through perspective projection; and
- the surface is manipulated to effect the reorganization of image details

In EPS the presentation transformation is separated into two steps: surface manipulation and perspective projection.

Our intention is to develop EPS to support exploration of presentation possibilities with regards to addressing screen real estate issues. The major decisions behind EPS endeavour to keep it general in scope and to address comprehension issues such as user disorientation when faced with the shifting spatial organization that results from the use of elastic presentation.

3.1.1 Developing EPS as a General Framework

To keep EPS general it has been developed independently from an application. This involves recognizing the distinction between representation and presentation and defining presentation space. Chapter 1 describes this distinction in detail. To re-iterate briefly, a representation is inherently tied to the characteristics of the information because creating a representation involves making a mapping from information, while a presentation of a representation organizes such things as the point of view and the relative emphasis of different parts or regions of the representation. Interactive presentation supports visual exploration of the representation.

Therefore, to work in presentation space we assume that there exists a representation for which visual exploration would be useful. When working directly with a particular information representation, one common design principle is to both respect and, when possible, take advantage of the information representation characteristics. To ensure generality we take precautions in order to avoid tying the framework to the current information's characteristics. This requires using a meta-representation whose characteristics are common to all representations that the framework is intended to include.

What are the basic visual representation characteristics? Even though the information that is being represented may well be multi-dimensional, and its corresponding visual representation may have skillfully incorporated many of these dimensions (see [13, 160, 161, 162]), the representation will need spatial coordinates in order to create a display. This spatially dimensionality is usually limited to one, two or three dimensions because of the nature of displays. We initially examine presentation possibilities for two-dimensional representations. By this limitation we are only assuming that the representation lies in a plane and can therefore be placed on a flat surface.

The decision to develop EPS for visual representations that have two spatial dimensions was made because: there are many two-dimensional visual representations, there is considerable information about creation and interpretation of two-dimensional representations, and most presentation research to date is intended for two-dimensional representations. Also, extrapolation to higher dimension visual representations is possible (Chapter 7). As any two-dimensional representation can be placed on a surface, the latter can be used as a meta-representation. Creating methods for presenting a surface will provide the generality necessary to handle all two-dimensional representations.

Furthermore, many of the methods that are tied to particular applications make use of a DOI function that depends on domain knowledge. In fact, it has been said that a detail-in-context method that does not use a DOI cannot be useful [118]. In contrast we believe that it is possible to create a framework that is independent of domain knowledge allowing us to think about distortion viewing in general. However, we do agree that when creating an information and task specific tool it is important to respect the particular information's characteristics [92, 166].

3.1.2 Developing Comprehensible Presentations

While the possibilities of elastic presentation seem to hold much promise in creating more readily usable visual exploration environments, we recognize that the distortion used to create them may be problematic as it distorts the information representation. If full advantage is to be taken of the elastic potential of computational presentation space, this comprehension issued will have to be faced. To this end we investigate methods of informing users about the presentation operations. This meta-knowledge is incorporated to convince the user that their visual exploration is a view operation and as such not damaging to the information.

Handling Context

The manner in which context is preserved has considerable effect on the resulting presentations (see Chapter 2). Detail-in-context research started from two basic ideas about maintaining context:, full compressed context [147], and filtered or sufficient context [52]. This, combined with the insights gained from observations [70, 137] graphical interpretations of the ideas now strive for full, if distorted, contexts [20, 84, 137, 138]. Consequently, we start by developing a method that will provide full context, considering the issues of spatial constancy, the mental map, and visual gestalt significant. Not until the introduction of *folding* (Section 4.3) is the possible advantages of sufficient contexts discussed. Furthermore, for the purposes of improved comprehension, we introduce the idea of maintaining as much of the context undisturbed as possible (Section 4.1).

Using the Surface

The representation of the information is considered to be separate from the two-dimensional plane or surface that contains it. This decision was reached from two reasons: to ensure generality and for the possibility of providing interpretation support. Placing the representation on a surface keeps the framework general because a solution that will work for a surface will be applicable to all two-dimensional representations.

In regards to interpretation support, distortions can be quite readable when applied to some representations such as regular layouts (particularly grids or text). Unfortunately, not all information representations can be arranged so regularly. It is necessary to be able to provide interpretation support for irregular and sparse layouts. The distinction between the representation and the surface on which it is located allows visual cues to be provided about the surface, creating distortions that are still readable even when there are gaps in the layout (see Chapter 6).

Using Three Dimensions

The initial decision to leave the prevailing 2D-to-2D approach and manipulate the surface in three-dimensions (before projecting back to two-dimensions) was taken primarily for comprehension reasons. The use of a three-dimensional distortion approach offers several advantages:

- there is greater freedom in manipulation in three-dimensions;
- the use of the third dimension provides a useful metaphor for the actions performed to create the distortions. Pulling a section towards oneself allows it to be better seen, or in this case magnifies it, which appears to be a natural response; and
- displaying the three-dimensional form of the pliable surface creates the possibility of comprehensible distortions because common understanding of three-dimensional shape supports interpretation of which sections are magnified and which are compressed with an intuitive notion of the extent.

Also, it interests us to explore the 2D/3D nature of a computational display. A computer can be used to display either 2D or 3D representations, however, the 3D representations resolve into 2D displays when held stationary. While this inherent characteristic has been discussed as problematic [98] it seems possible that it may extend some advantages. As a step towards exploring these possible advantages, we reverse this situation. While the 2D detail-in-context presentation is the primary interest, with motion an EPS presentation resolves to 3D.

In its static form the computer offers a similar display medium to printed graphics. However, one of the computer's major differences from printed graphics is that the images displayed can change over time. As the computer's 3D display space is created and manipulated mathematically, computational ability to change displayed images over time extends into many possibilities.

Perspective

As the manipulations will be performed in 3D and a computer screen is 2D a some type of projection is needed. We use perspective projection to view the manipulated surface because its visual effect most closely resembles human visual perception. We benefit from the use of perspective geometry and from the common ability to read depth in perspective images. Even though the question of whether understanding three dimensions from perspective is an innate ability or a learned skill has not been definitively answered [57], the ability to read perspective can probably be assumed to be common in computer users.

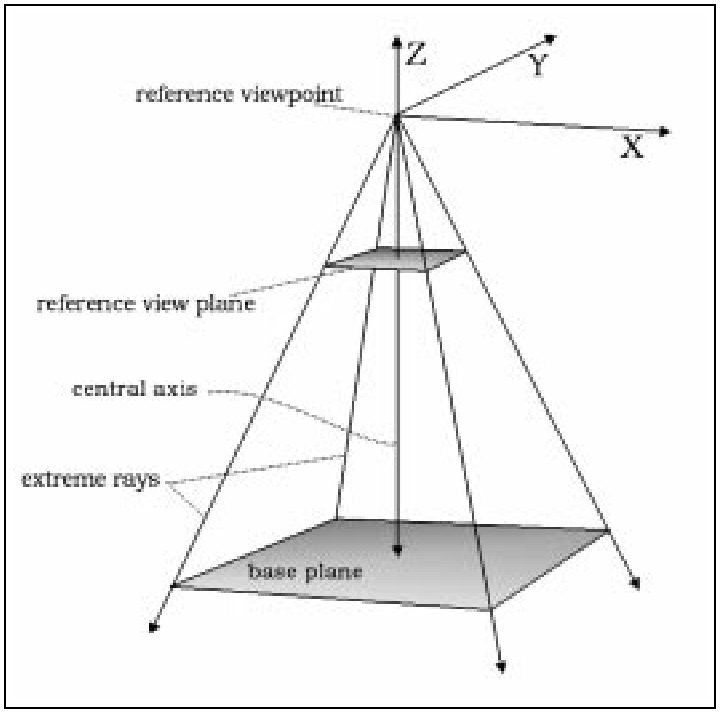

Figure 3.1: The reference view point

Perspective viewing is used in two different ways. One, the usual three-dimensional viewing which allows for changes in viewpoint and provides rotational viewing of the three-dimensional space and the surface within in it. Two, for the purposes of creating a detail-in-context presentation, one particular perspective view, the *reference view*, is considered. Presentations are created with respect to a fixed viewpoint and a fixed view plane. These locations are called the *reference viewpoint* (RVP), and *reference view plane*.

RVP is located at the origin, looking down the negative z axis. The reference view plane is centred on the z axis with its 2D axes parallel to x and y, as is the *base plane*. Placing the surface on the base plane is taken be the *normal* presentation; that is, the two-dimensional representation is not considered to be either magnified or compressed when it is positioned flat on the base plane.

Locating RVP on the z axis, above the base plane provides single point perspective views for objects aligned with the x, y plane. Since geometric relationships of angles, parallelism, and proximity are preserved on all planes that stay parallel to the reference view plane, maintaining the orientation of an x, y plane during any translation preserves these relationships. Therefore x, y planes that are translated in z, change in scale only. An increase in z magnifies them, and a decrease in z uniformly compresses them. This action corresponds to zooming. Also translating the x, y plane in x or y corresponds to panning and scrolling respectively. That this simple explanation of panning, scrolling and zooming exists within EPS has allowed us to consider how these features relate to the detail-in-context presentations that can be developed within EPS.

It is with respect to the RVP that detail-in-context presentations are developed. The distinction between RVP and the system viewpoint allows viewing of the 3D manipulated surface used to create the detail-in-context presentations from other angles such as to the side or below. While the system viewpoint changes through rotation, once established, RVP remains fixed. Figure 3.1 is an image taken from a viewpoint that provides a view of the RVP and the *reference view volume*.

3.2 Developing a Detail-in-Context Approach with EPS

As discussed in Chapter 2, one frequently identified computational presentation problem is the conflict between presenting sufficient detail and maintaining adequate context. As a first test of the usability of EPS, we use it to develop a detail-in-context method, the Three-Dimensional Pliable Surface (3DPS) [20]. For EPS to serve as a general presentation framework it must be possible to use it to develop a detail-in-context presentation method that does not have caveats. That is, this method should be able to provide, at the very least, all of the features that have already been discussed in the literature as desirable. A variety of methods have been suggested that address the problem of providing presentations that support detail-in-context readings. These methods have focused on individual applications, creating results that are tied to the information that is being presented and the tasks being performed. In spite of this, there is a growing consensus in the literature about which features a detail-in-context approach would ideally have. Listed below are the features that have been repeatedly discussed in the literature as desirable in a detail-in-context presentation method.

Desired Features

1. **Providing a Focal Point** A detail-in-context presentation needs at least one focal point that is capable of displaying the required detail while still located within its context. This is usually achieved by magnifying the focus and compressing the context sufficiently to provide the extra display space needed by the magnified focus. Ideally the context is compressed such

that none of it is visually eliminated. This basic requirement is achieved, albeit differently, by all detail-in-context methods.

2. **Choice of Focal Shape** To create a focus the user should be able to specify the shape and size of the area of interest according to their information and their current task. Not all regions of interest will have the same shape. Most previous methods offered either point or rectangular foci. The closest approximation to arbitrary focal shapes was the inclusion of convex polygons as foci [138].

3. **Control of Focal Magnification** Ideally a user would be able to control both the degree of magnification or detail in the focal region and whether or not the scaling is preserved in the focal region. In many situations it is desirable that the magnified focal regions are scaled only. While distortion is often used to create detail-in-context views, it can be important in some tasks (such as reading distances on a map) that focal regions are kept clear of distortion. While no other approach supports precise specification of the degree of focal magnification, Magnification Fields [85] uses an iterative method to approximate degree of magnification.

4. **Freedom of Focal Location** The user should be able to specify the location of the focus according to their information and current task. This freedom is supported in other 2D-to-2D approaches, however, it appeared that freedom of focal location was only possible with a single focus if the approach used three dimensions and perspective projection [132].

5. **Multiple Foci.** There are many situations where having only a single focus is too limiting. Often a task will require more than one focus. Many approaches support multiple foci in some manner, however, usually with caveats such as causing ghost foci, placing limitations on focal proximity, and adjusting previously set focal sizes. The earliest, if non-computational, detail-in-context method Polyfocal Display [80] provides uncompromised multiple foci, and subsequent to our work Non-linear Magnification [84, 85] and Focus Line [59] also support uncompromised multiple foci.

EPS uses a two step process to create presentations. A two-dimensional representation is placed on a surface. The surface is manipulated in three dimensions and then viewed in two dimensions through use of perspective projection. This is a 2D-to-3D-to-2D process. This chapter describes the use of EPS to fulfill these listed features in the development of our prototype Three-Dimensional Pliable Surfaces (3DPS) [20]. Due to our decision to use a 3D approach, it is not immediately apparent that this will be possible. For instance, at the onset of this research, it appeared that a three-dimensional approach and multiple foci were mutually exclusive![132]. Also, it was not until after our work first appeared that the methods based 2D-to-2D transformations provided all suggested functionality in this lists [59, 84, 85].

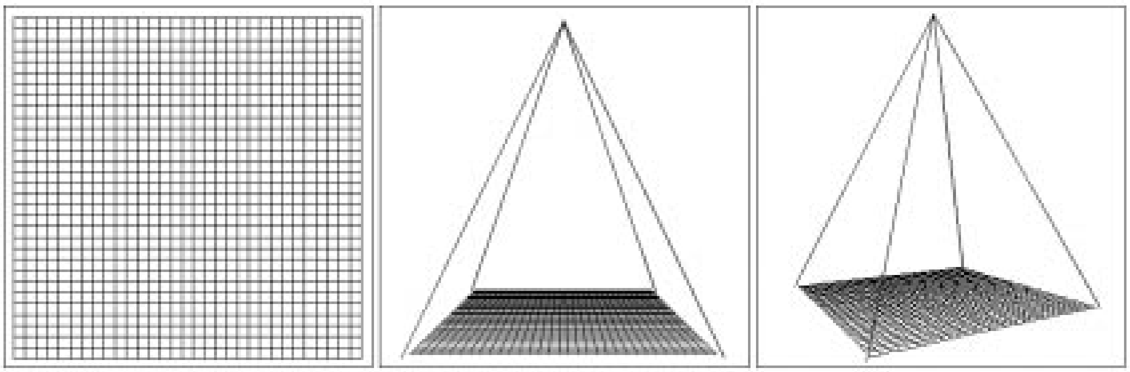

Figure 3.2: A regular grid has been placed on the surface which is in normal position: on the left, projected view from RVP; centre and right, side views

This explanation is illustrated with figures created with our prototype, 3DPS [20]. To make the surface manipulations easy to read a regular grid is placed on the surface. Figure 3.2 shows three perspective views of the regular grid on the flat surface in its normal position. The left-hand image of Figure 3.2 shows the view from the RVP when the RVP and the projection viewpoint are coincident. The centre and right images are side views showing the flat surface in normal position, the RVP, and the extreme rays of the reference view volume. Side views and/or top (from RVP) views are used as appropriate throughout this explanation to illustrate both the three-dimensional shape of the surface as well as the projected view. The projected view from the RVP will show the detail-in-context presentations.

3.3 Providing a Single Focal Point

Creating a detail-in-context presentation requires revealing sufficient detail of a chosen focus while the focus is still located within its context. Focal detail is revealed by magnifying the focus. The space required for the magnified focus is achieved by compressing the context. We start by explaining how to use 3D surface manipulation and perspective projection to create a single central magnified focus with sufficient surrounding compression.

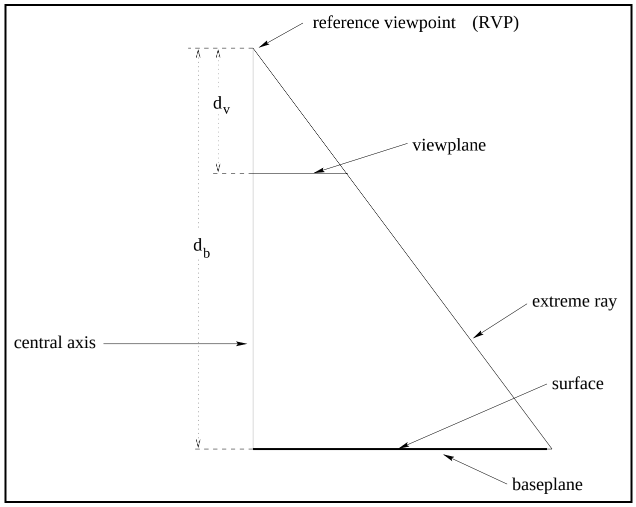

Figure 3.3: A cross-section diagram of half the view volume from RVP to the base plane and from the central axis to the extreme ray with the surface in normal position

When the surface is in its normal position on the base plane it is considered to have no magnification (Figure 3.3). At the top of the triangle is the RVP and at the bottom is the base plane. At the left of the triangle is the central axis and on the right of the triangle is an extreme ray or the edge of the viewing volume. The surface is in normal position on the base plane. Figure 3.2 shows perspective views of the start position.

In EPS, magnification is a function of distance in z from the surface to the RVP. Translating the surface in z towards RVP lessens this distance and increases the magnification. We describe the

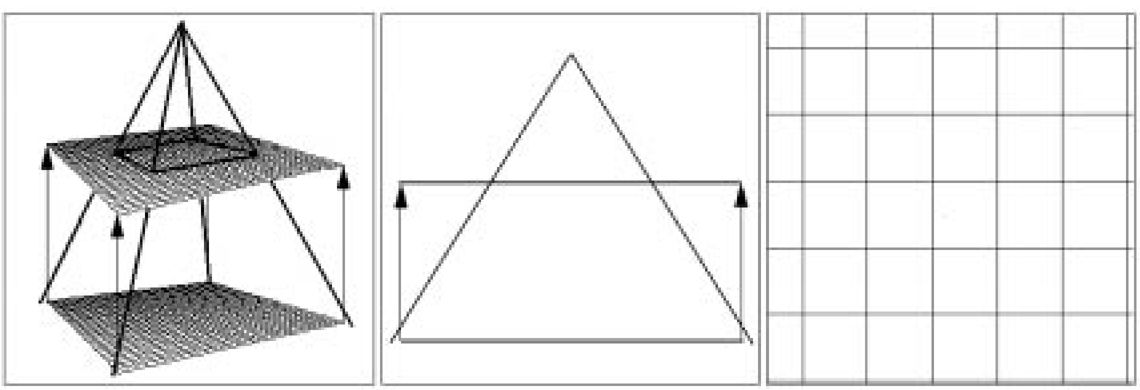

Figure 3.4: Magnifying the grid with simple zoom: left, 3D side view; centre, cross-section view; right, resulting magnified view showing the centre of the magnified grid

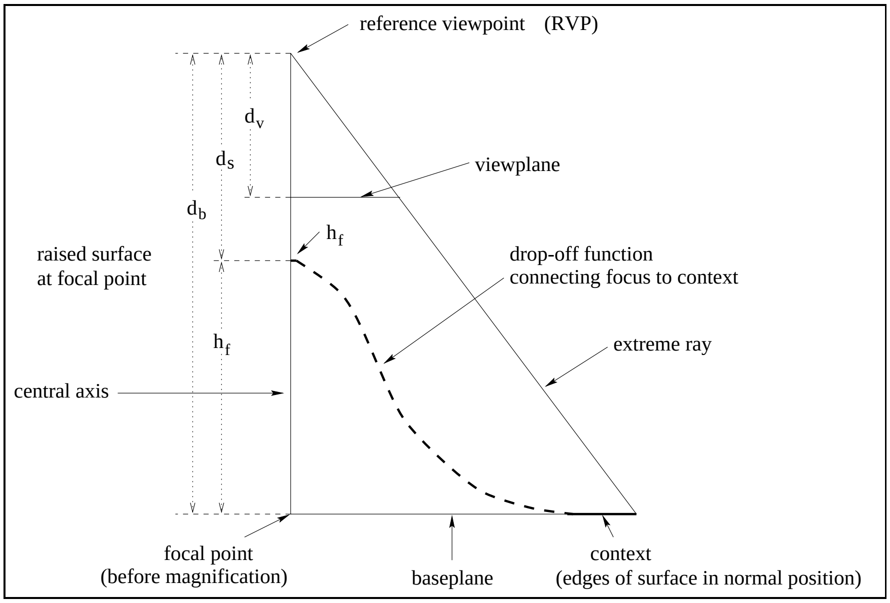

Figure 3.5: A cross-section diagram showing the raised central focal point, the profile of the drop-off function and the edge of the surface in normal position

basic procedure for obtaining detail-in-context before discussing the precise details about specifying the exact degree of magnification for the focal region in Section 3.5. If the entire surface is raised uniformly with the central focus, the effect is that of zooming with its attendant problems of loss of context. Figure 3.4 shows that only the portion within the view volume is visible as magnification is increased. To achieve detail-in-context, we want to raise the surface at the focal point and keep the edges of the surface on the base plane within the view volume. Raising the focus effectively magnifies it, but as the magnified focus occupies more screen space, the context will have to be compressed. The combination of magnification and compression required to create a detail-in-context presentation is achieved by raising only the specified focus to the height that will provide the magnification and using a drop-off function to position the rest of the surface. Figure 3.5 shows the surface at the central focal point translated in z to a height h_f and the edge of the surface still in normal position on the base plane. h_f is the distance in z from the base plane the focus. A function specifying the drop-off from the focal height h_f to the context on the base plane is needed.

The compression is caused by both the slope of surface and the distance z that the surface is from the RVP. Choice of the drop-off function is crucial as it will affect both the visual result and the performance. To maintain full context with smooth visual integration and minimal distortion, a Gaussian drop-off function is used. Figure 3.7 shows a side view of a manipulated surface for a

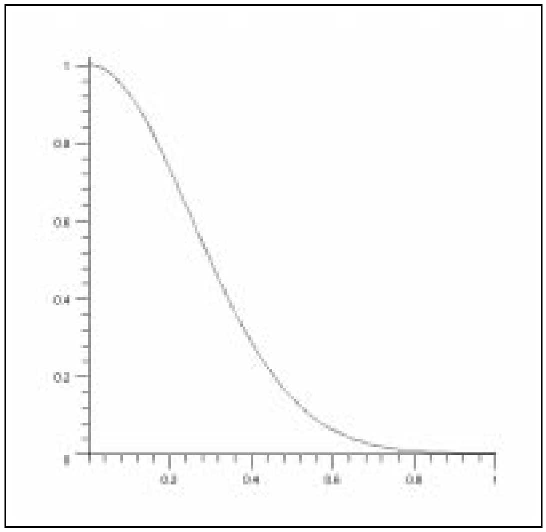

Figure 3.6: The Gaussian drop-off function

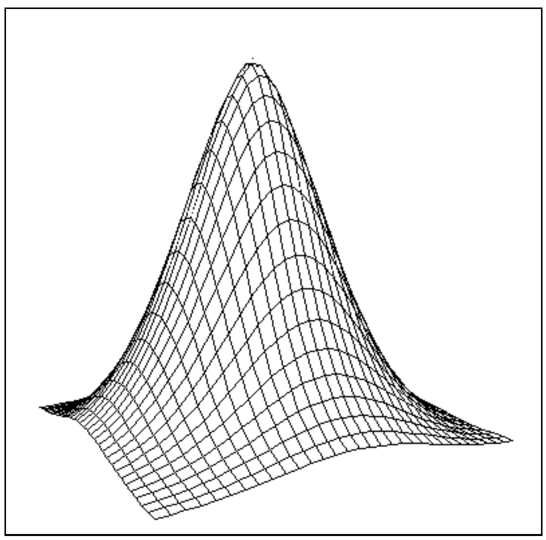

Figure 3.7: A single point focus with a Gaussian drop-off function

single central focus. The centre point of the focus is raised and the rest of the surface follows the profile of a Gaussian curve. Visually, the Gaussian curve is a natural choice as its bell shape curves away from its apex and inflects to curve gently back to the base plane. This provides smooth visual integration from the focus into the compressed region of the context and again from the compressed region into the untouched context surrounding it. Furthermore, it is a simple mathematical function with no discontinuities in curvature. We use a steep Gaussian with the standard deviation $\sigma = 0.1$

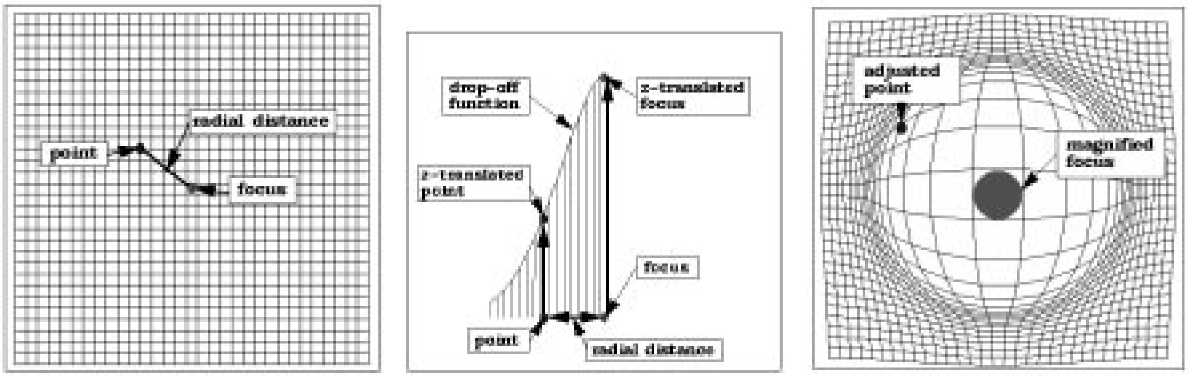

(a) The focal point is selected

(b) A half cross-section view, showing the curve of the Gaussian drop-off function

(c) The detail-in-context presentation with a magnified central focus

Figure 3.8: The process of creating a single central focus.

to ensure that it has no perceptible height when it reaches the edge of the context (Figure 3.6).

Figure 3.8 illustrates the process used to establish a detail-in-context view. First the focal point f_c is selected; in this case it is a central focus (Figure 3.8(a)). Then, for each point p on the surface, the radial distance d_p from f_c is calculated. The focal point f_c is raised perpendicularly from the base plane, directly towards the RVP, to the chosen height h_f (Figure 3.8(b)). This height h_f is the apex of the Gaussian curve and is set to produce the desired focal magnification. To create the compression, each surrounding point p is set at a height h_p according to the Gaussian drop-off function:

$$h_p = h_f \cdot \exp^{-(\frac{(d_p)^2}{\sigma})}$$

where d_p is the 2D radial distance from the focal point f_c to p. Figure 3.8(c) shows the detail-in-context presentation. The focus f_c has been magnified and the point p has been displaced radially and slightly magnified.

In summary, the manipulation of the surface provides magnification and compression in the appropriate regions. In Figure 3.9, the leftmost image shows the side view of the three-dimensional Gaussian curve in the centre of the field of view. The centre image in Figure 3.9 shows a cross-section view revealing the perpendicular translation vectors throughout the Gaussian curve. The rightmost image shows magnification/compression pattern for a single focus in the centre of the field of view.

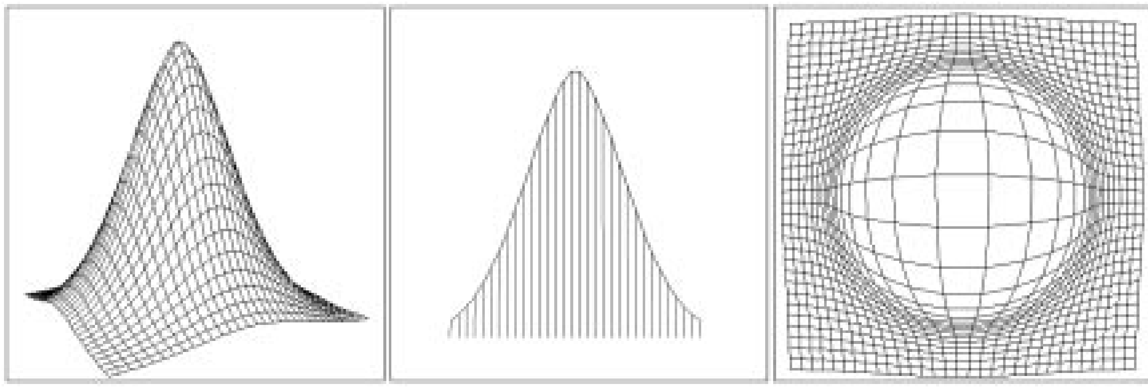

Figure 3.9: Single focus at the centre of the field of view: left, 3D side view; centre, cross-section view showing projection vectors; right, resulting detail-in-context view

3.4 Freedom of Focal Shape

With a point focus, as just described in Section 3.3, magnification is provided at the centre of the focus and then the degree of magnification immediately begins to taper off (Figure 3.9). This creates some degree of distortion right across the focus. There are many situations where this is not acceptable and one would like an entire region that is magnified to scale. For example, when presenting MRI images to a radiologist, (detail-in-context presentations can be effective for an array of images where a single image is a focal region), it is very important for diagnostic purposes that there is no distortion within a single image [166]. Also, as data exists in great variety with many types of characteristics, it is important to provide a diversity of focal types. In fact, an ideal focus could well be a region with a fairly non-standard shape such as a river valley or a political boundary. Focal regions of arbitrary shape allow a user to specify a focus to suit their current needs. Point, circular, line, poly-line, concave and convex polygonal foci can be created within our framework.

Achieving a point focus is described in Section 3.3 and illustrated in Figure 3.10(a). To create a circular focus (Figure 3.10(b)) for a specified focal radius f_r, all points whose distance d_p to the focal centre is less than f_r are translated to the focal height h_f. For calculating the height h_p for points outside the focal region, distance d_p that is used with the drop-off function is taken as being from the point p to the edge of the focal region.

The procedure for creating point foci can also be extended to include line and poly-line foci. In the case of line foci (Figure 3.10(c)), for each point p on the surface the shortest distance d_p from p to the focal line is calculated. This distance is then used to calculated the point's height h_p from the drop-off function in the same manner as previously discussed. All points on the line will be

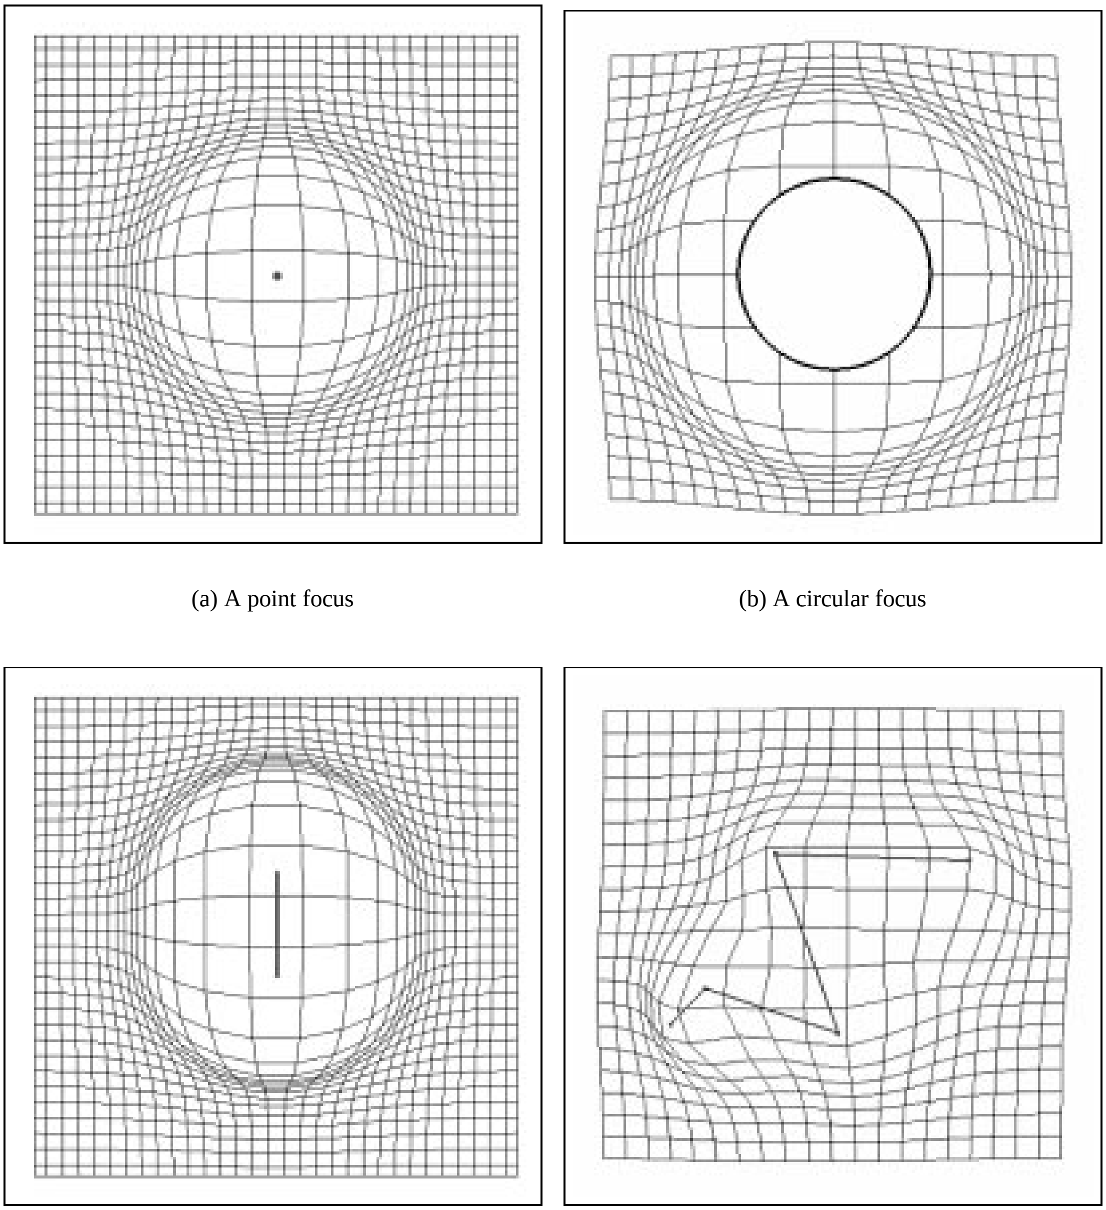

(a) A point focus

(b) A circular focus

(c) A line focus

(d) A polyline focus

Figure 3.10: Focal types

projected to the requested height h_f (Figure 3.10(c)). In this manner a poly-line can be used as a focus (Figure 3.10(d)).

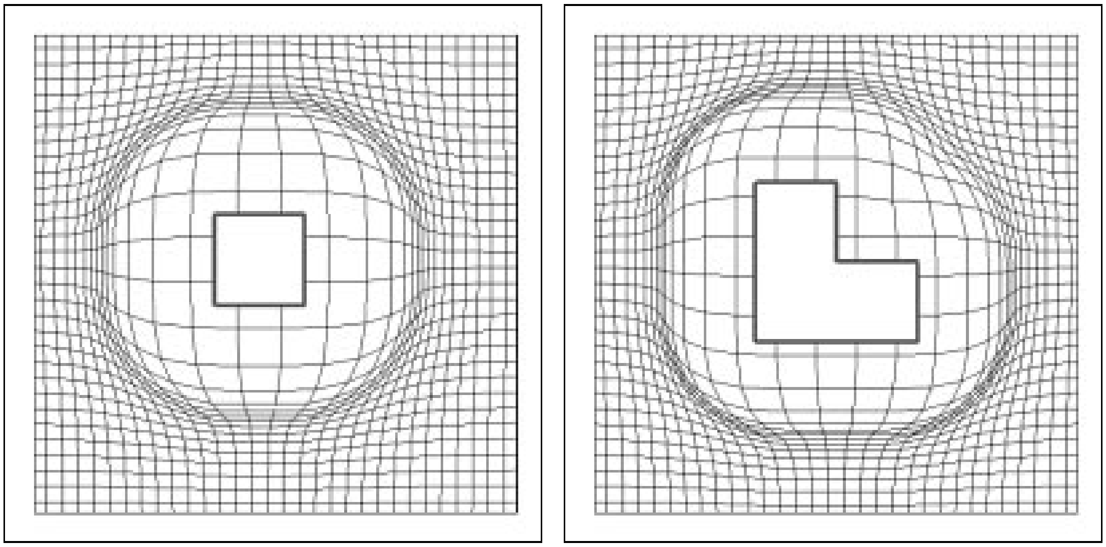

(a) A convex focus

(b) A concave focus

Figure 3.11: Focal types

Arbitrary polygon foci are also possible. A point in polygon test is done before calculating the distance d_p. If the point is either on the line or within the polygonal focus, it is projected to the full height h_f. All other points are positioned by finding the closest distance to a polygon edge and applying the drop-off function to calculate h_p. The height h_p of a point outside of a focus but within an affected region is determined not by its distance from the center of the focal region but by its distance to the edge of the defined focus. Figures 3.11(a) shows a square focal region and Figure 3.11(b) shows a simple concave polygonal focus.

In this framework providing arbitrary focal shapes is not difficult. In general, for each point p: if p is inside the focus, it is placed at focal height h_f, if p is outside the focus then find the distance d_p to the closest edge, the height h_p is calculated from d_p, with drop-off function. However, the point in polygon test slows the algorithm, as now the complexity is also dependent on the number of polygon edges.

We consider the focus to be the point or region that is selected by the user. A *lens* includes the focus and its compensating region of distortion. A lens is *global* if its effects extend throughout the representation to the edge of the context. A lens is *constrained* if the effects of the lens are limited to a sub-region of the representation.

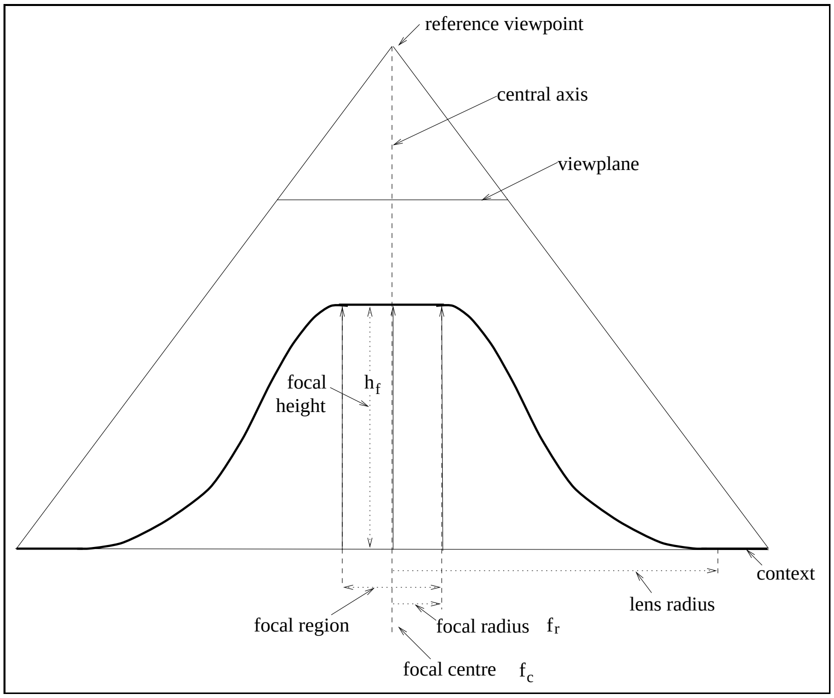

Figure 3.12: The parts of a lens

3.5 Magnification Control

A major issue is user control of the degree of magnification at the focus. In EPS, magnification can be readily controlled by adjusting the focal height h_f. The series in Figure 3.13 shows, from left to right, the effect of increasing the height. Users are able to increase the magnification until they

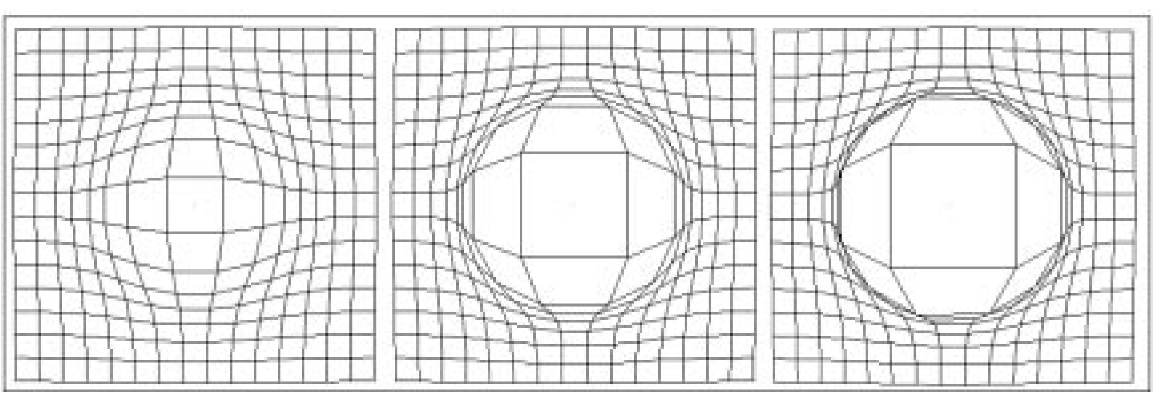

Figure 3.13: Single foci: effects of varying height only

feel satisfied with the amount of detail that they can see. In many cases this is sufficient and with all previous detail-in-context approaches a method of subjective judgment was the most magnification control that was provided. However, it can be important to know the degree of magnification of the focus and to be able to specify it precisely.

Also, the amount of magnification created in response to a translation in height h_f increases as the distance to RVP decreases. In fact, this increase becomes so dramatic as RVP is approached that the user does not have fine control of the level of detail. Figure 3.14 shows the magnification

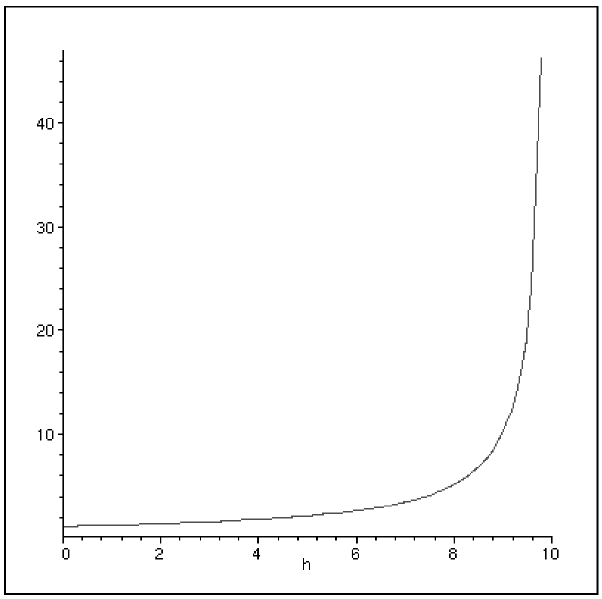

Figure 3.14: The magnification function with RVP to base plane distance d_b of 10. Note the rapid increase after the distance $(d_b - h_f)$ is less than one.

response as a function of height h_f. As RVP is approached, a tiny increase in h_f will cause a huge increase in magnification. A previous solution, used in Document Lens [132], substituted $\log(h_f)$ for h_f. This helped to provide user control at higher degrees of magnification but still did not allow the user to specify the degree of magnification. Other previous methods also have provided the user with inexact methods of controlling magnification. For example, the focus may resize itself in response to duration of mouse contact [7]. As long as there is no need for precise control or knowledge of the degree of magnification these methods may be adequate. Our method provides fine user control with precise specification of the degree of magnification.

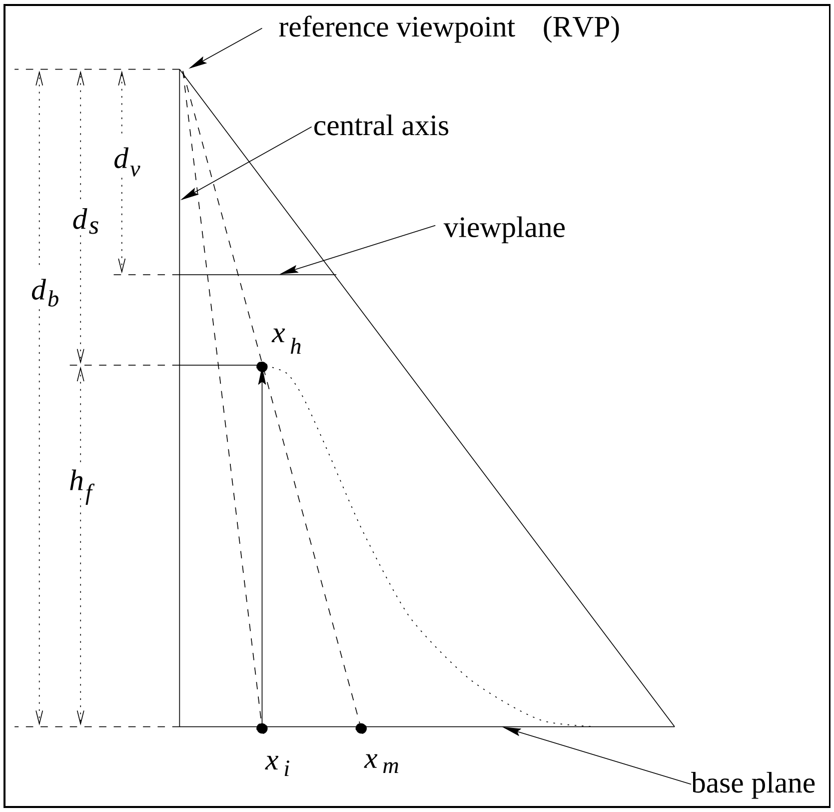

Figure 3.15: The process of manipulating the surface

Since the magnification function of 2D-to-2D methods is the derivative of their transformation function, it can be difficult to obtain a transformation function given magnification information [94, 85]. Magnification Fields [85] uses an iterative approach to approximate a presentation based on requested magnification. EPS relies on the basic geometry of a perspective view volume, the properties of which provide a single step conversion between magnification and transformation. This combination simplifies the mathematics of the relationship between magnification and transformation to the geometry of similar triangles.

There are three factors in EPS; the actual translation in z, the perceived magnification and the apparent translation in x and y. During the surface manipulation, focal points or regions are translated in z to a height h_f and all other points p within the lens are translated in z to a height h_p according to the drop-off function. When viewed in perspective from RVP, the focus appears magnified. However, as is evident in the side views of the manipulated surface the focal areas are not actually expanded. They appear magnified because of perspective foreshortening. Translating a portion of the surface to a height h_f in z creates a change in perceived magnification mag. Though points have only been translated in z with no x and y movement, points within the lens appear to have also been translated in x and y, when viewed in perspective from RVP. Perspective projection appears to move a point from its initial location at (x_i, y_i) to (x_m, y_m).

RVP is centred directly above the surface along the z-axis, at an established z distance, d_b. Both the apparent transformation in x and y and the resulting magnification factor mag can be readily calculated from the surface height h_f. Additionally, the height can be derived given either the magnification factor mag or the x and y translations. Figure 3.15 illustrates this for x_i.

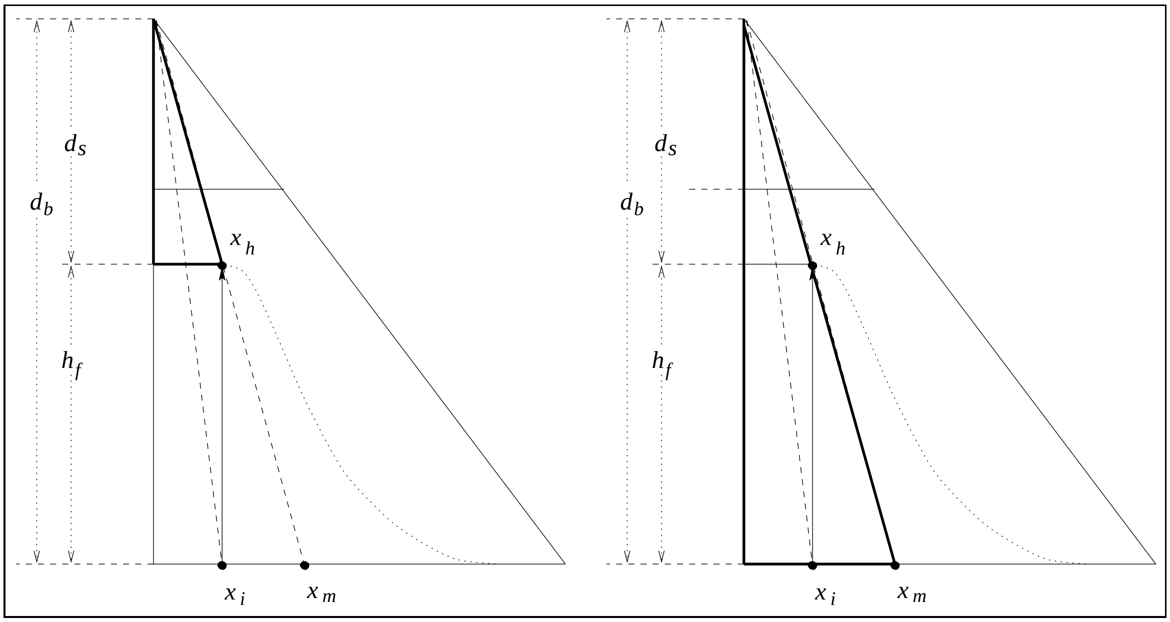

Figure 3.16: Similar triangles show the relationships between z-translation and magnification. The lateral distance between the projections of x_i and x_h on the view plane is the perceived magnification.

From similar triangles (see Figure 3.16)

$$\frac{(x_m)}{(d_b)} = \frac{(x_i)}{(d_s)} \qquad and \qquad \frac{(y_m)}{(d_b)} = \frac{(y_i)}{(d_s)} \tag{3.1}$$

Given: Surface Height.

A point $p_i(x_i, y_i, z_i)$ on the base plane (Figure 3.16) is translated a z distance h_f to locate it at $p_h(x_i, y_i, (z_i + h))$. The projection of p_h on the view plane is the same point that x_m, on the base plane, would project.

Since d_s is known and given h_f, the coordinates x_m and y_m can be calculated as follows:

$$d_s = d_b - h_f \tag{3.2}$$

$$x_m = x_i \cdot \frac{(d_b)}{(d_s)} \qquad and \qquad y_m = y_i \cdot \frac{(d_b)}{(d_s)} \tag{3.3}$$

The magnification factor is:

$$mag = \frac{d_b}{d_s} \tag{3.4}$$

These coordinates allow the option of performing transformations directly by translating the point in x and y or through perspective by adjusting its height. The x and y translation corresponds directly to other transformations that operate in the plane [84, 108, 137] and is discussed further in Chapter 5.

Given: Magnification Factor.

Given a magnification factor mag, the z displacement h_f is obtained from d_b.

The magnification factor is:

$$d_s = \frac{d_b}{mag} \tag{3.5}$$

$$h_f = (d_b - d_s) \tag{3.6}$$

The coordinates x_m and y_m are calculated as in Equation 3.3.

Given: Transformed coordinates.

In the projected view, the transformed image is the result of apparent lateral translations x_m and y_m. If the known quantities are the transformed coordinates x_m and y_m of the apparent lateral translations, the initial coordinates x_i and y_i can be obtained with h_f as follows:

$$x_i = x_m \cdot \frac{(d_b - h_f)}{d_b} \qquad and \qquad y_i = y_m \cdot \frac{(d_b - h_f)}{d_b} \tag{3.7}$$

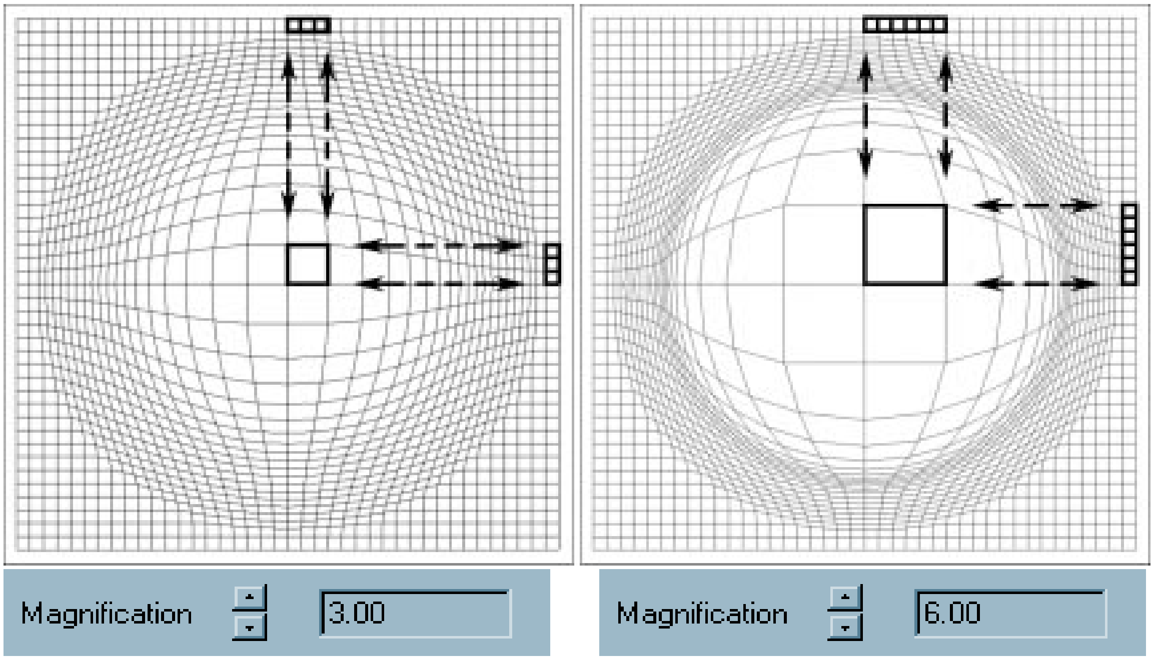

Figure 3.17: Magnification of a uniform grid by 3 times and 6 times respectively.

The magnification factor is the same as noted above. To obtain either the the magnification factor or the initial coordinates, both the transformed coordinates and the height h_f are needed.

The simplicity of the relationships between magnification and displacement is one of the major advantages of basing our framework on perspective projection. We use the magnification factor to compute the height to which the surface should be maximally displaced, and the drop-off function interpolates the magnification smoothly into the context where there is unit magnification. This allows direct specification of focal magnification. Figure 3.17 shows magnification of a grid by three and six times respectively.

3.5.1 Scaled-Only Focal Regions

A change in scale is a change in the degree of magnification (or compression) with no accompanying distortion. Two aspects to user control of scale at the focus are: the choice over whether or not there is a scaled-only region at the focus and how the degree of magnification can be controlled. There are two simple methods for creating scaled only regions at the focus. One method is by selecting a region, whether it be a polygon or circle, and placing all points within it at the focal height. The other is by limiting the lens' maximum height. While both of these methods provide scaled-only focal regions the visual integration is quite different.

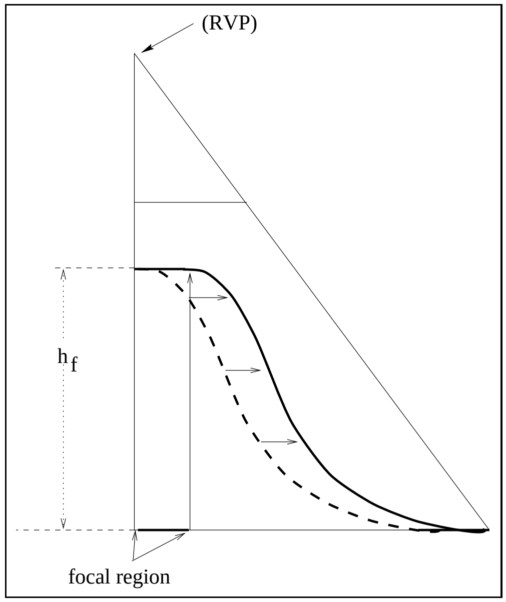

Figure 3.18: Creating a region of uniform scale by expanding the focal region

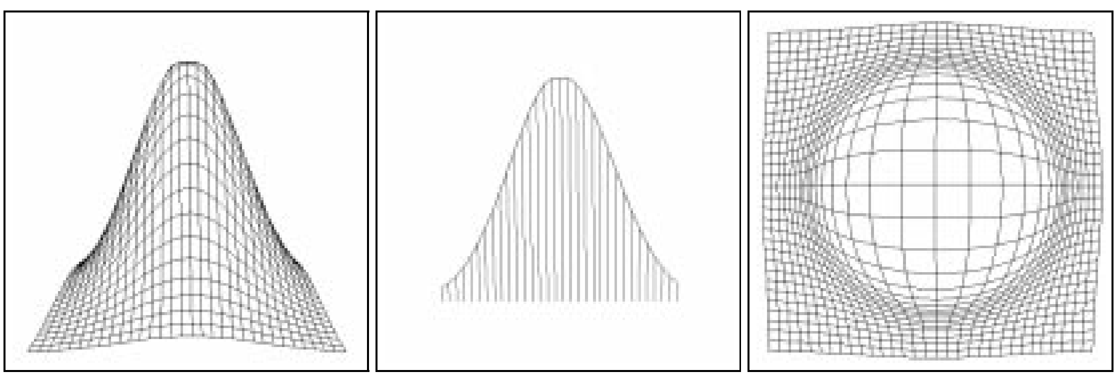

Figure 3.19: Creating a scaled-only focal region by expanding the focal radius

With a point focus, the bell shape of the Gaussian curve provides a smooth connection between the focus and its immediate surrounds, magnifying adjacent regions almost as much as the focus. This type of visual integration is maintained when the focal region is expanded, and the effect of the drop-off function is started at the edge of the focal region. Here the maximum lens height h_m is the focal height h_f and the 2D radial distance used for the Gaussian drop-off function is $(d_p - f_r)$. Figure 3.19 has a scaled-only region that is two squares wide. The profile of the side and vector views show how the Gaussian curve starts from the edge of the scaled-only region, curving gently

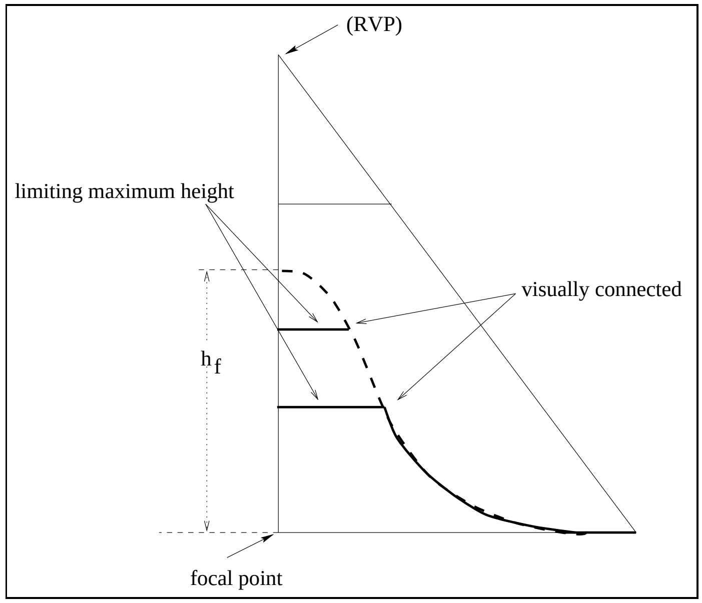

Figure 3.20: Creating a region of uniform scale by limiting the lens height

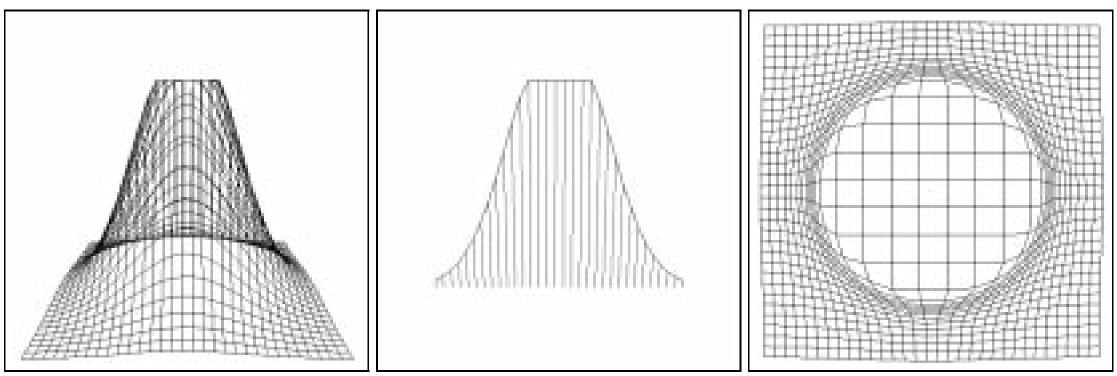

Figure 3.21: Creating a scaled-only focal region by limiting the lens height

out from the focal region. As in any visually integrated distortion the edges of scaled-only focal area blend into the context.

Creating a scaled-only region can also be done by limiting the the drop-off function at a set height h_m. Specifically, the center of the 3D Gaussian curve can be projected up to the height h_f, which may be higher than h_m. To provide a flat region where only scaling occurs, the lens may be truncated or limited to h_m (Figure 3.21). The points of the surface in the central magnified region are all projected up to the same height h_m. A point p is set to $\max(h_p, h_m)$, where height h_p is calculated from h_f, d_p, and the drop-off function as before. This method has the effect of cutting straight through the lens. As shown in Figure 3.21 there is a change in direction of the surface at the edge of the scaled-only region. This makes a more definite break in the visual integration and places the more extreme compression directly around the scaled-only focus. Figure 3.22 shows a larger scaled-only region, with simple visual connection providing information about where scaled-only areas end and distorted areas start.

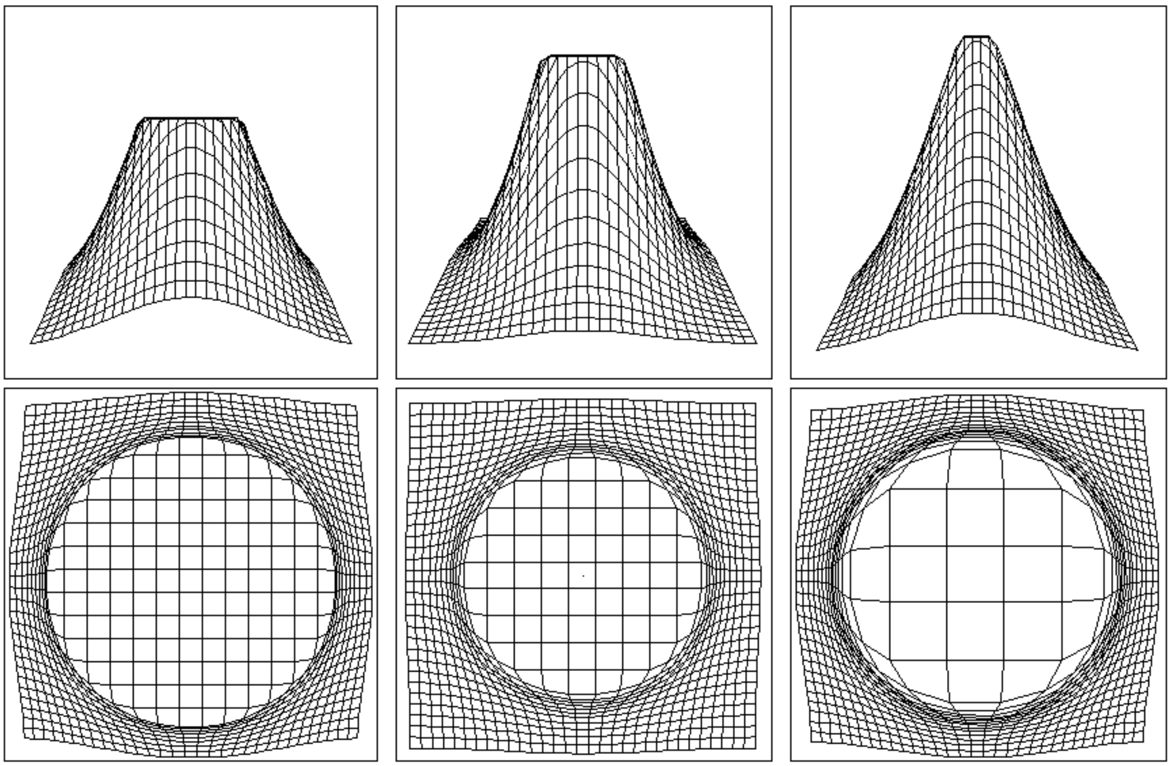

Figure 3.22: Three different single foci with flattened tops. The top row shows the profile view of the truncated curve and the bottom row shows the matching projected view. The grid clearly shows the regions of uniform scale. Note also the change in scale corresponding to curve height

3.6 Freedom of Focal Location

The solution just described for creating a single central focus does not function well in other locations. Direct application of the perpendicular translation towards the reference viewpoint works at the centre of the field of view. In any other location, changes in focal height can easily place the section of interest outside the view volume, where it may be magnified but can no longer be seen. If

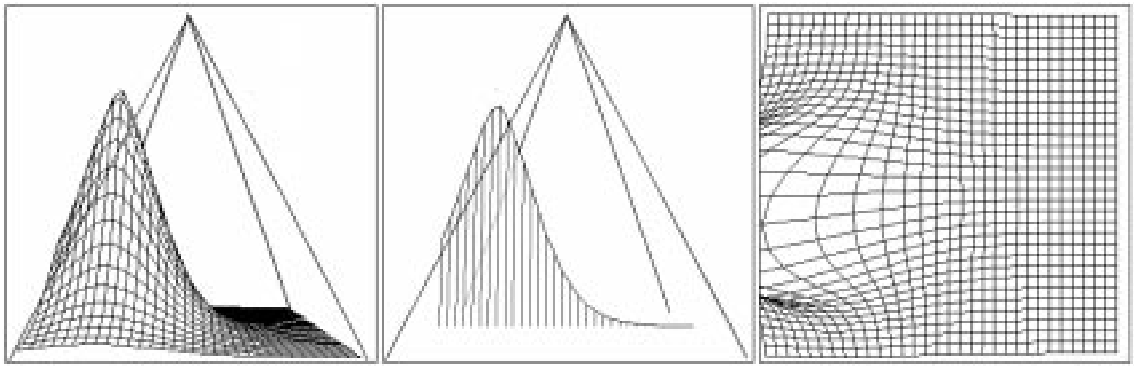

Figure 3.23: Single off-centre focal point: left, side view of the off-centre perpendicular focus; centre, the cross section view showing perpendicular vectors and the view volume; right, the projected view showing that the focal region cannot be seen

the region of interest is located at a point other than that directly below the viewpoint on the original surface, projection perpendicular to the plane of the surface causes the focal region to move out of the viewing volume as it increases in magnification. This is illustrated in Figure 3.23 where the leftmost image shows the side view of a single perpendicular focus in conjunction with the view volume. Though magnified, the actual region of interest is outside of the view volume. Figure 3.23, centre, shows the perpendicular vectors and the image Figure 3.23, right, shows the resulting projected view. This is a fundamental geometric limitation of the configuration of the viewing volume used in perspective projection.

As the intention is to allow focal points to be selected from any location this situation must

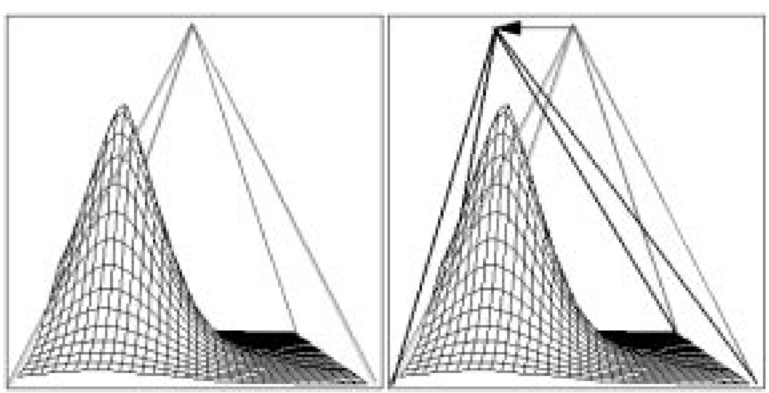

Figure 3.24: Single roving focus solution: left, side view of perpendicular focus; right, the side view of perpendicular focus with the translated viewpoint

be resolved. Mackinlay et al. [99] and Lamping et al. [91] limited their approaches by freezing a single focus at the centre of the field of view and moving the information across it. Robertson and Mackinlay [132] provide freedom of location for a single focus by translating the viewpoint to keep it directly above the focus (Figure 3.24). Both of these solutions are limited to a single focus. We provide freedom of focal location without eliminating the possibility of multiple foci.

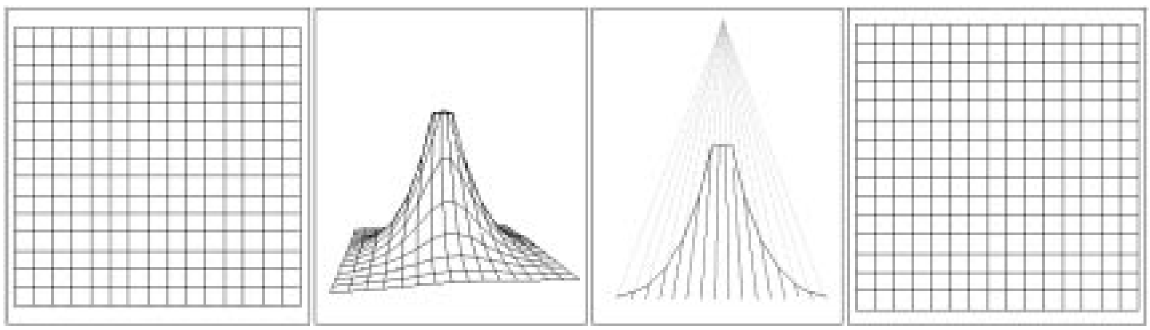

Figure 3.25: No magnification series: left, flat untouched grid; centre left, 3D side view of actual shape of adjusted grid; centre right, side cross-section view showing the vectors to the reference viewpoint; right, projected view of the adjusted grid showing no evident change

The translation vectors in our central focus are perpendicular to the base plane and point straight at the reference viewpoint (Figure 3.9). If the translation vectors for a focus in other locations also pointed at the reference viewpoint they would stay in the view volume and therefore still be visible. However, directing all vectors towards the reference viewpoint has no resulting magnification effect. This is because as all these vectors converge towards the viewpoint, the amount of magnification achieved by the change in height exactly corresponds to the amount of compression achieved by their convergence. That is, if a point $P(x, y, z)$ is moved to any position along the vector from its current location to the viewpoint, there is no perceptual change in the projected image. This results in the situation where the surface can take on any type of formation when seen from the side but the projected view will always be the same.

Figure 3.25 illustrates this with four images. The leftmost image shows the top view of the uniform grid before any manipulation. The centre-left image shows the three-dimensional side view, the centre-right image shows the cross-section view, revealing the vectors all pointing towards the reference viewpoint. Both these centre images show that the surface has been manipulated, but the rightmost image, which shows the projected view of the transformation in the centre two images, appears exactly as the untransformed image. The applied compression through repositioning in x and y is exactly compensated for by the change in z.

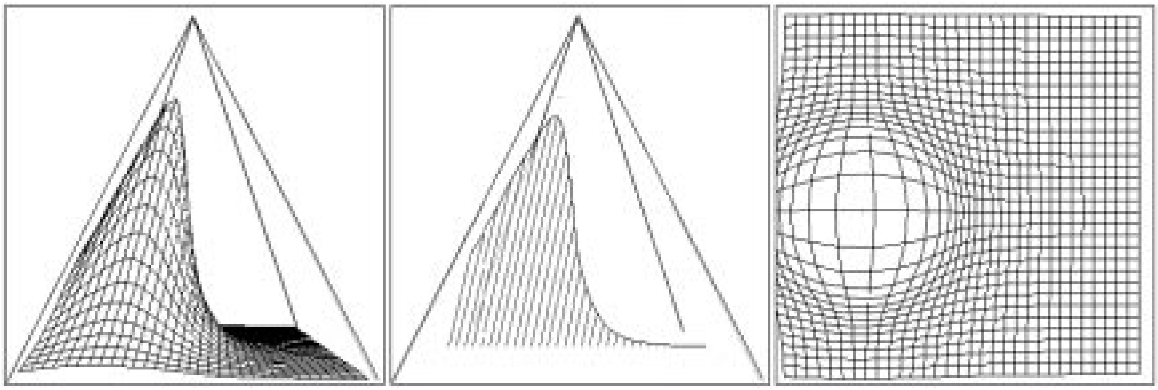

Figure 3.26: Viewer Aligned Focus: left; 3D side view showing the focus directed to the viewpoint and inside the view volume, centre; side profile view showing parallel vectors with the central vector directed to the viewpoint, right; shows the top view of the projected image

The geometry in Figure 3.9 reveals that the distance from the reference viewpoint combined with the parallel translation of the surface creates the magnification effect. We create a *viewer-aligned focus* by directing the translation vector for the centre of the focus to the reference viewpoint and ensuring that all other translation vectors for this focus are parallel to the central *viewer-aligned vector*. The viewer-aligned vector directs the focus towards the reference viewpoint, keeping the focus within the reference view volume. Translating all points in the focal region along vectors that are parallel to the focus' viewer-aligned vector creates the magnification. Figure 3.26 shows a viewer-aligned focus, the side view shows that the focus is inside the reference view volume, the cross section view shows the parallel vectors, and the right image shows the top view of the projected result.

Viewer-alignment allows any focal point to be selected. However, with traditional normalization of the viewer-aligned vectors the magnification response will vary depending on the point chosen. The height or z translation of the unit viewer-aligned vectors is a function of its position on the base plane. The further a viewer-aligned vector is from the centre of the base plane, the more acute its angle is to the base plane. Therefore, viewer-aligned vectors at the edge of the base plane will have less z translation and correspondingly less magnification than those in the centre. The important factor is not the length of each viewer-aligned vector but that the z-components of these vectors be the same. Therefore the z-component of the viewer-aligned vectors is normalized to one instead of the length of the view vector. Figure 3.27 illustrates this, showing how the traditionally normalized vectors (left) describe a gentle curve while the z-normalized vectors (right) form a straight line.

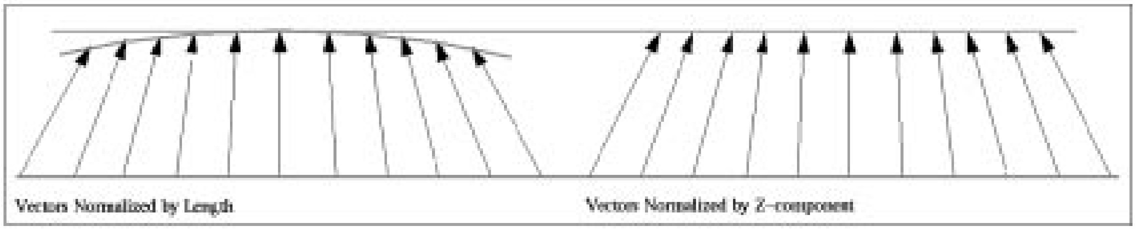

Figure 3.27: On the left, traditionally normalized vectors and on the right, z-normalized vectors

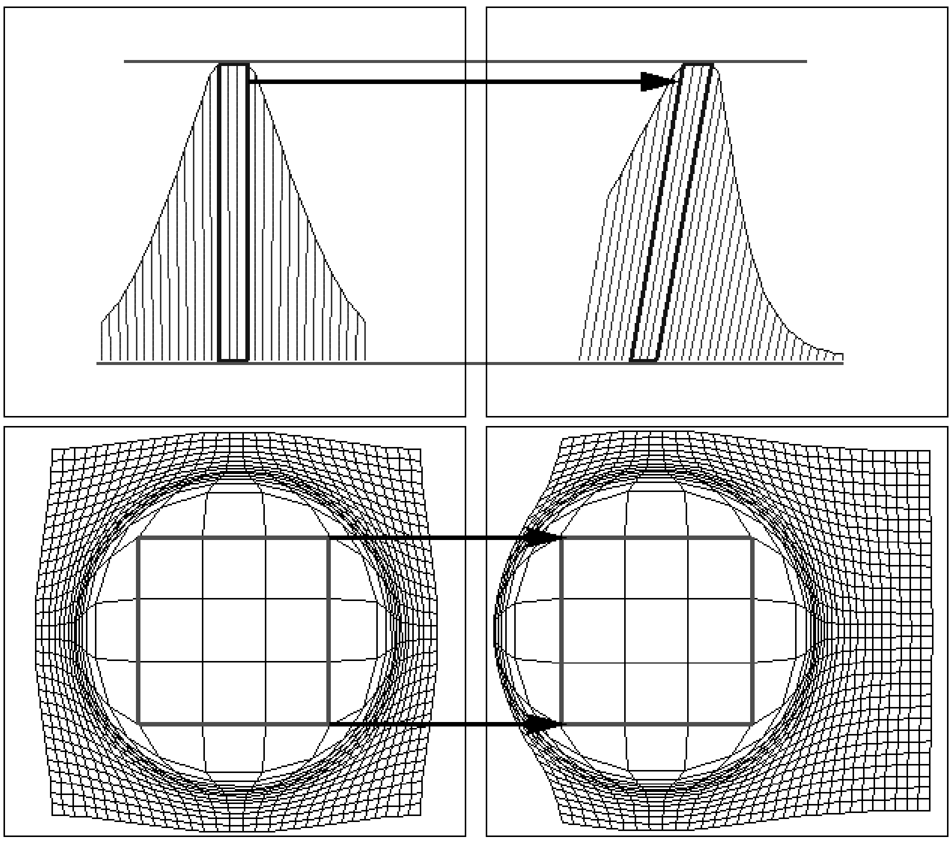

Figure 3.28: Scaled-only viewer-aligned focus

Using z-normalized vectors ensures uniform magnification response across the entire view.

The discussion about providing focal regions that are scaled-only considered those foci that are located at the centre of the field of view. The use of viewer-aligned vectors for arbitrarily located foci does not affect this ability to provide regions of uniform magnification. Figure 3.28 illustrates this. The cross-section vector images in the top row show the rectangle formed from the focal region's position on the base plane, the translation vectors, and the focal region's translated position. The effect of the z-normalized viewer-aligned vectors is to sheer this rectangle. The size and height of the top does not change. Therefore the scaled-only focal magnification will stay the same as the focus is moved across the base plane. The projected views are shown in the bottom row of Figure 3.28. For arbitrarily shaped foci the center of the arbitrary region is used to determine the viewer-aligned vector.

3.7 Multiple Foci

The viewer-aligned solution lends itself naturally to multiple foci since it makes it possible to locate foci at will. However, another problem arises. When two focal points are close to each other there can be a *buckling* or visual discontinuity where they meet. Figure 3.29 shows the regular grid buckling between two foci, both from the cross-section vector view on the left and the projected view on the right. If the information is a vector representation or point and line layout such as a graph, this buckling will result in information reversal. For instance, in a map the east/west ordering of two cities might be reversed. If the information representation is raster or an opaque surface there will be a crease and some region of the information representation will not be visible. This problem is not unique to this approach; Sarkar et al. acknowledged that their multi-focal approach [138] based on morphing [12] caused this information reversal when two large foci were relatively close.

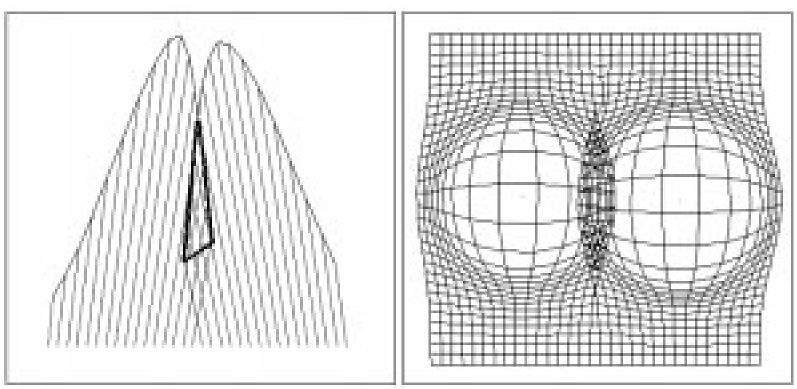

Figure 3.29: Buckling: left, the cross section vector view, clearly showing the surface buckle; right, the top or projected view showing the surface buckle

Each foci has a centre, a translation height, a viewer-aligned vector, and a curved region of gradually decreasing influence that is defined by the drop-off function. A given point on the original

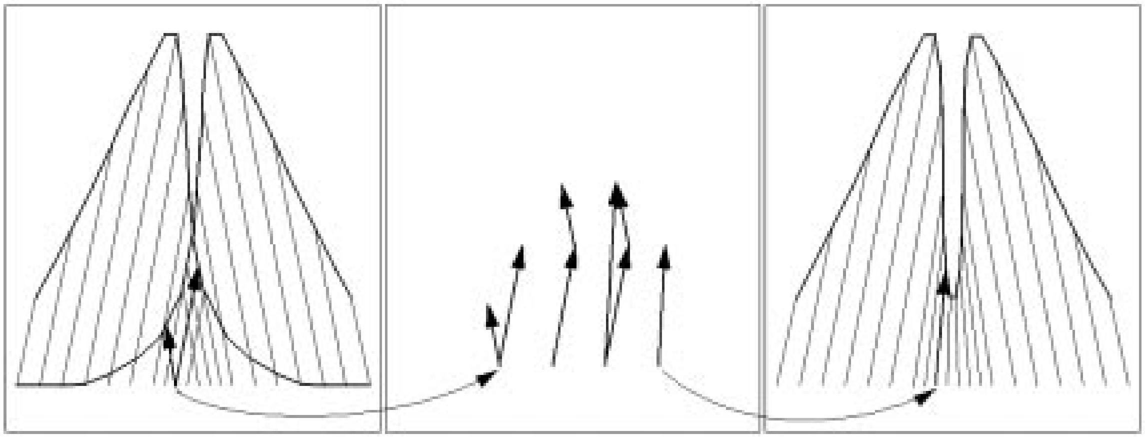

Figure 3.30: The blending process

surface that lies between two foci and inside both of their regions of influence could be projected according to the drop-off function from either focus. The dominant focus is the one with the greatest height at the point's location and should be the one seen when viewing from the reference viewpoint. Therefore the point is projected according to the dominant curve. For the curves of any two adjacent focal regions there is a point where this dominance shifts from one curve to the other. If the viewer-aligned vectors differ, which they will since the two foci are just adjacent not coincident, there will be a buckle. When the height dominance switches from one curve to the next the difference in projection vectors can cause a reversal in the direction of the surface, producing the surface loop seen in Figure 3.29.

The opportunity to examine exactly what occurs in these regions was provided by initially writing our prototype 3DPS to visualize the algorithm. Figure 3.29, right image, shows the top view of two foci colliding. The open mesh clearly shows the buckled region but is not very explanatory as to why this is happening. Figure 3.29, left image, shows a cross section view of the translation vectors. Here the nature of the buckle is more explicit. The surface loop is created as the slope from one focus slides under the other. As two or more foci move close to each other, the angles of their viewer-aligned vectors differ. By maintaining parallel projection throughout a given focus the angle of the viewer-aligned vector is propagated to the edges of the distortion of the lens. Thus where two foci meet there can be a sharp change in the angle between the translation vectors.

Our solution keeps the translation vectors viewer-aligned in the focal regions and adjusts their alignment in the inter-focal regions so that there is no abrupt change in translation direction. This is achieved by using a weighted averaging of the viewer-aligned vectors as follows:

For every point to be projected
  For every curve that affects this point; (drop-off function is non-zero)
    Calculate the point's height for each curve as if it was projected perpendicular,
      to the baseplane
    Store the highest one, as this is the dominant height at this location and will be the
      final translated height.
    Obtain each curve's viewer-aligned vector (calculated earlier from the curve's
      center and RVP and then z-normalized).
    Calculate the translated vector for each curve (the translated height times the
      viewer-aligned projection vector).
    Add all the point's translation vectors together obtaining a total translation vector.
    Z-normalize this total translation vector to obtain the blended translation vector.
    Use the saved dominant height and the blended translation vector to calculate the
      point's actual position.

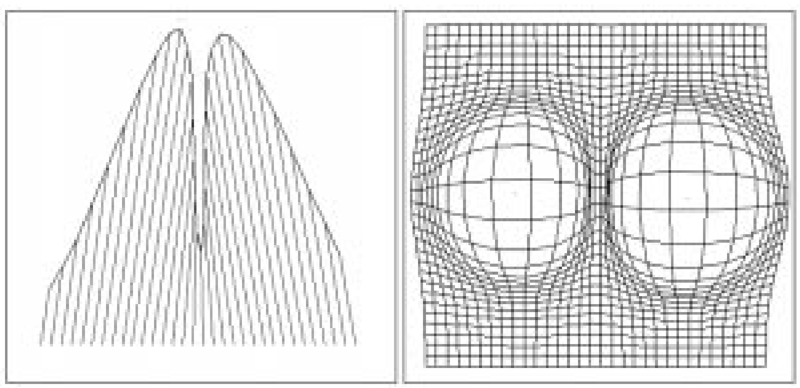

Figure 3.31: A blended inter-focal region

Figure 3.30 illustrates this process. In the left image, for one point in the inter-focal region the translation vectors for both curves are highlighted. The centre image shows the blending. The two contributing translation vectors are added and then z-normalized. The resulting translation vector uses the blended z-normalized vector and the dominant height. Figure 3.31 shows a profile view of the blended translation vectors and a top view of the blended surface. This vector blending solution is simple enough to be readily interactive and works well with any number of focal points.

A consideration is that blending allows a curve that extends right under another curve's focal point to affect the distortion at that focal point. While the actual height is maintained the change in vector tilt will have some affect, possibly causing a focal region that had been scaled only to now be distorted. This is not an absolute solution, as it is still possible with unconstrained freedom of slope to create very extreme situations where buckling will occur. This is to be expected as it is possible to cause information reversal with curve slope alone.

3.8 Discussion

We have introduced a distortion-based appraoch, 3-Dimensional Pliable Surface (3DPS) [20], which allows for multiple arbitrarily-shaped foci on a surface that can be manipulated by the viewer to control the level of detail contained within each region.

While the simple mathematical relationships based on perspective projection, as described in Section 3.5 are easy to use, it is important to also consider computational complexity issues. Multi-scale views of discrete layouts are created by adjusting the individual points. If the transformation is applied to a grid, the grid points can be used to adjust a surface which can hold any 2D information or image. As each point in the grid must be visited, the size and resolution of the grid is the factor that primarily affects the complexity of the algorithm. The grid sizes we have been using (which allow interactive rates) vary from (10x10) to (30x30). Our method, and all other approaches with the exception of [85], are also dependent on the number of foci. At every grid point a calculation is made to determine the influence of each focus. Considering the literature that indicates that humans tend to hold seven plus or minus two static objects in memory [142], or Pylyshyn's studies [128] indicating that this number is closer to four plus or minus one for moving objects, in normal use the number of foci will rarely exceed ten. This means that for each point in the grid we generally will not exceed ten focal operations.

In contrast, the direct use of magnification fields by Keahey et al. [85] causes iteration through the grid for each point. The number of these iterations is influenced by the degree of magnification requested, with more iterations as magnification increases. The iteration continues until the achieved magnification is within some error tolerance.

If the algorithm can be reversed mathematically the resulting view can be used in many ways other than simply seeing the detail situated in its context. These include the ability to edit and annotate distorted images. Furthermore, being unable to readily return to previous presentation layouts can increase the possibility of feeling lost in computer space.

For an information presentation to be useful it is essential that the information can be understood. If the distortions are read as part of the information this may lead to false interpretations [160]. It is important to provide support to allow one to perceive the distortions as distinct from the information.

The points regarding maintaining a unified presentation, recognition and interpretation are discussed in Chapter 6.

3.9 Verifying the Usefulness of EPS

EPS is a framework for understanding presentation possibilities for two-dimensional information representations that:

- separates development of presentation concepts from the needs of a particular application or type of information,
- is independent of an application in that it is not dependent on the characteristics of the application's information, or representation,
- provides control of several display parameters such as degree of magnification and location of maximum compression, providing new control possibilities,
- explains existing presentation methods, making both distinctions and relationships between them more accessible,
- allows extrapolation between existing methods, creating an understanding of a presentation space,
- opens up new presentation possibilities by allowing incorporation of features from what previously seemed to be incompatible methods into a single method,
- takes advice from cognitive science, endeavouring to create presentations that are not known to be cognitively strenuous and can utilize visual gestalt by retaining the perception of the information space as a single event,
- and has proven extensible to three-dimensional representations.

Having established that EPS can be used to develop a detail-in-context method, Chapter 4 describes using EPS to extend of the idea of user control of the amount of magnification versus the amount of compression and explains how the idea of re-positioning foci can be supported because of the basic structure of EPS. This is analogous to the freedom of re-positioning areas of interest that are displayed in separate windows. However, it is still within the unified elastic paradigm of varying the magnification within a single image. Also in Chapter 4, we discuss changing the distance metric.

One of the hallmarks of a useful framework is its ability to explain and/or relate previous research. Chapter 5 shows how EPS can relate previous methods ranging from full zoom ZUIs through magnified insets to detail-in-context distortion approaches. Furthermore, all the previous methods that can be explained with EPS can, if implemented through EPS, incorporate the full functionality as described in this Chapter and the extensions as described in Chapter 4, while retaining their characteristic patterns of magnification and compression. Furthermore, EPS supports extrapolation between these existing methods, creating both new methods and an understanding of elastic presentation space. Relating the separate methods algorithmically also allows inclusion of more than one method within a given application.

Chapter 6 offers a preliminary investigation of how to use those aspects of EPS that were included as comprehension aids. In developing EPS, we limited the dimensionality of the representation to two dimensions. In Chapter 7 we investigate removing this limitation.